Building an Open Science Grid

Ruth Pordes, Dane Skow Fermilab

representing the Open Science Grid Consortium

Grid2003
Shared Grid Infrastructure - 2004

-Goal to build a shared Grid infrastructure to support opportunistic use of resources for stakeholders.

Stakeholders are NSF, DOE sponsored Grid Projects (PPDG, GriPhyN, iVDGL) and US LHC

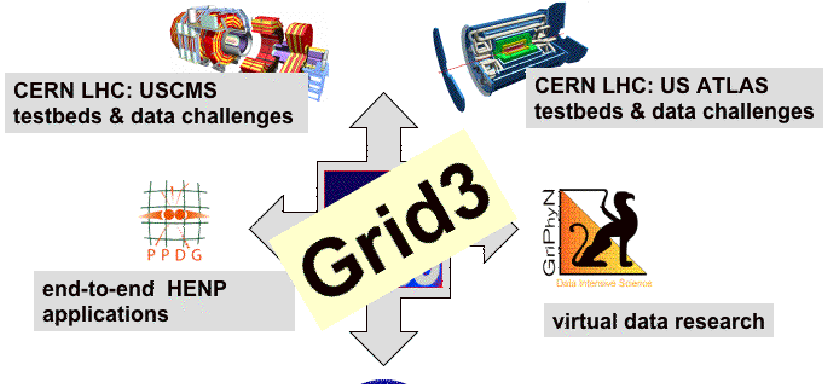

Grid3 Is a Success !

Team of computer and domain scientists deployed (simple) services in a Common infrastructure and interfaces across existing computing facilities.

Operating stably for over a year in support of computationally intensive applications.

Added communities without perturbation.

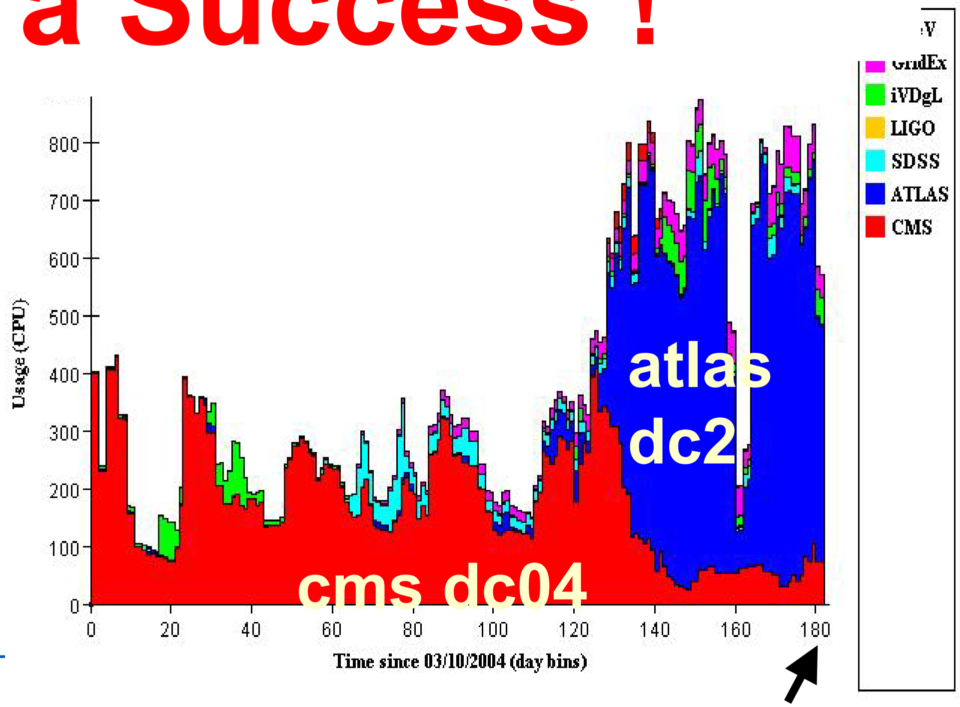

Grid Services Offered:

- Compute Element
  - Gateway through Globus GT2 GRAM; Support for 5 Batch systems
  - Minimal installation requirements on job execution nodes.
- Data Management
  - Data movement through GridFTP.
  - Space management through published disk areas ($APP, $DATA, $TMP)
- Workflow Management
  - Planning through GriPhyN VDS, Pegasus, VO specific schedulers.
  - Job Execution management through Condor-G, DAG, GridMonitor,
- Monitoring, Information & Accounting
  - Parallel systems for completeness: GT2 MDS, ACDC, MonaLISA, Ganglia, GridCAT
- User Authentication
  - LCG/EGEE Virtual Organization Management Service (VOMS)
- Operations
  - Grid Operations Center (iGOC)
  - Grid Testers: Exerciser, GridCat

Grid3 Resources Continue to Grow

Jan. 2004

New sites come through existing VOs or through agreement with Steering Committee.

Sites verification scripts test readiness.

Grid3 is resilient against new sites and applications and minor s/w upgrades.

Parallel “Grid3Dev” of ~7 sites used for new and updated service testing and verification.

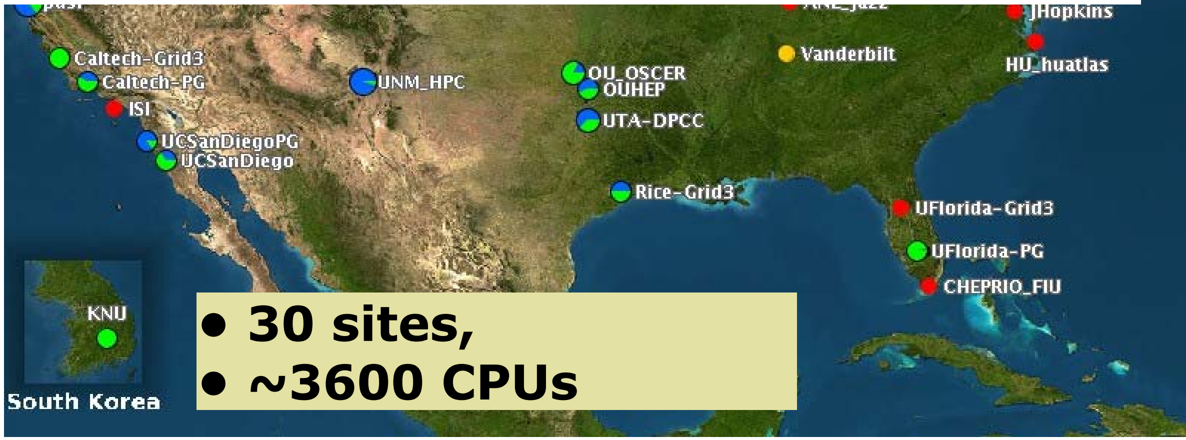

- **30 sites,**
- **~3600 CPUs**

Bioinformatics: Genomic Searches and Analysis

- Searches and find new genomes on public databases (eg. NCBI)
- Each genome composed of ~4k genes
- Each gene needs to be processed and characterized
  - Each gene handled by separate process
- Save results for future use
  - also: BLAST protein sequences

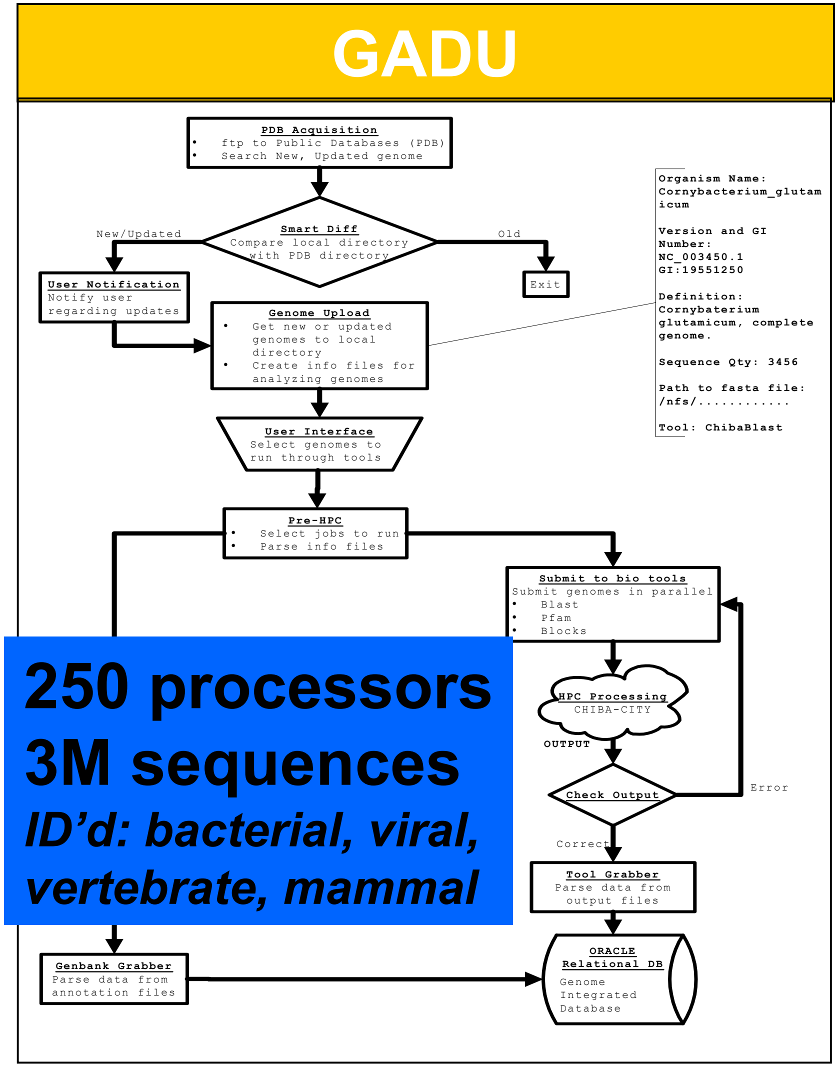

250 processors
3M sequences
ID'd: bacterial, viral, vertebrate, mammal

Astrophysics: SDSS Job Statistics on Grid3

Advanced Computational Data Center (ACDC) Job Monitoring

Grid3 Detailed Job Analysis

Temporal summary starting date: May 1 2004

to ending date: September 1 2004 inclusive

for: All resources

VOs share sites with simple priorities established through the batch system.

71949

Total CPU Time:
774 CPU Days

Average Job Runtime:
0.26 Hr

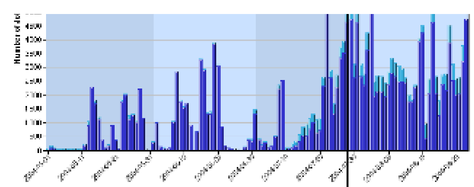
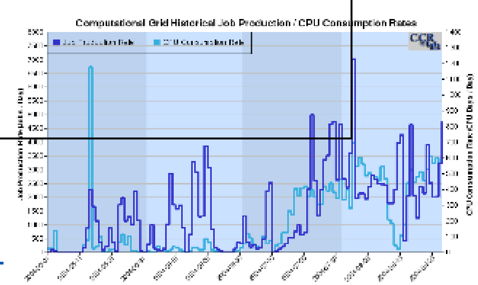

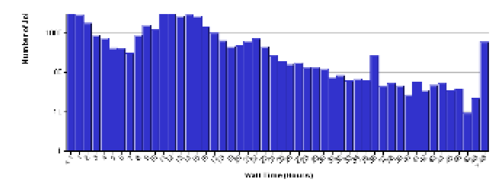
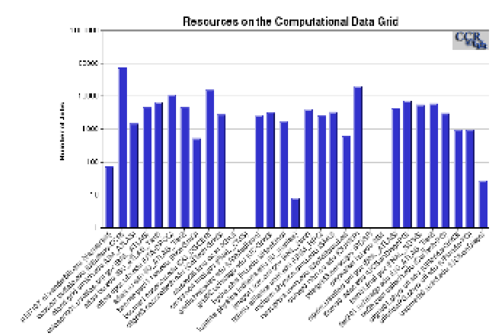

Open Science Grid

Open Science Grid
A Multi-Disciplinary
Sustained Production Grid

- Grid built and maintained as a coherent consistent infrastructure from
-

Adiabatic Evolution of Grid3 !

- Shared and opportunistic use of resources for executing jobs from all contributors.
- Open to science contributors.
- Partnering with other Grids for interoperability and coherency.
- Inclusive of small sites and organizations and usable as a Computer Science Laboratory.

OSG seeded by the US LHC

Open Science Grid

- LHC experiments and in particular US LHC software and computing committed to critical path reliance on Production Grids data analysis.
- Building system to manage and provide access to
  - <7PB distributed storage by2008
  - <3MSpecInts computation by 2009
  - ~8 Regional Centers distributed globally serving ~100 University distributed globally to serve ~2000 physicists.
- US LHC will present its resources to the Open Science Grid and actively contribute common services and validation of the infrastructure.

Facilities support Application Community Grid Environments through Common Interfaces and Infrastructure

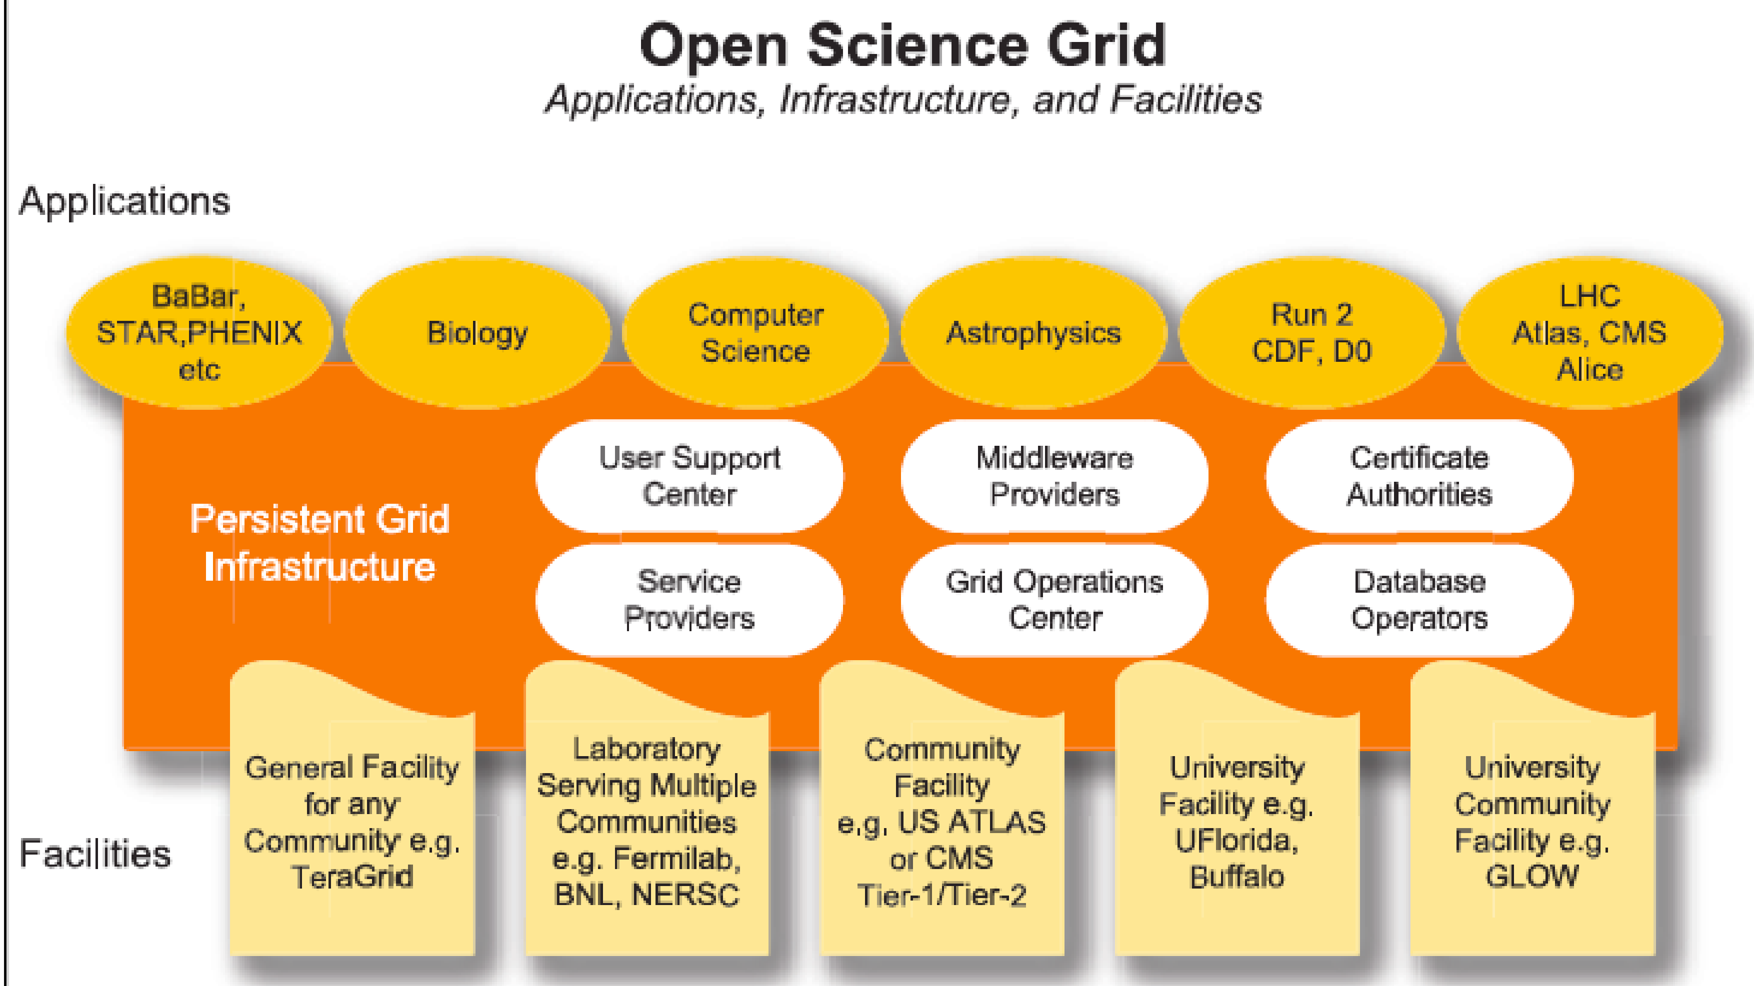

Open Science Grid

Character of Open Science Grid

- Distributed ownership of resources with diverse local policies, priorities, and capabilities.
- Guaranteed and opportunistic use of resources provided through Facility<-> VO contracts.
- Validated supported core services based on Virtual Data and NMI Toolkits. (currently GT2)
- Adiabatic evolution to increase scale and complexity.
- Services and applications contributed from external projects. Low threshold to contributions and new services.

OSG Organization Structure

Open Science Grid

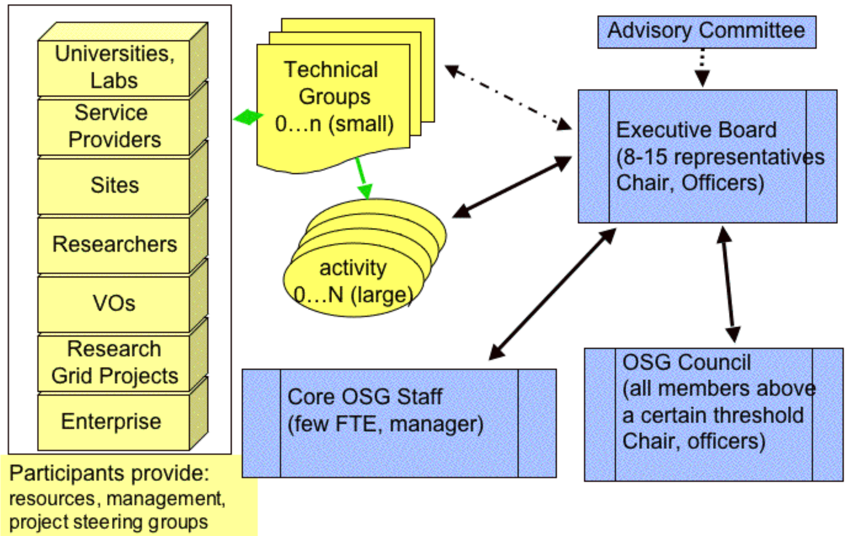

Activities:

- Integration
- Deployment
- Security Incident Response
- SE Service Readiness
- Site Account Mapping Service Readiness
- Discovery Service Readiness
- Operations

Technical Groups:

- Security
- Storage
- Education
- Monitoring & Information
- Policy
- Support Centers
- Governance

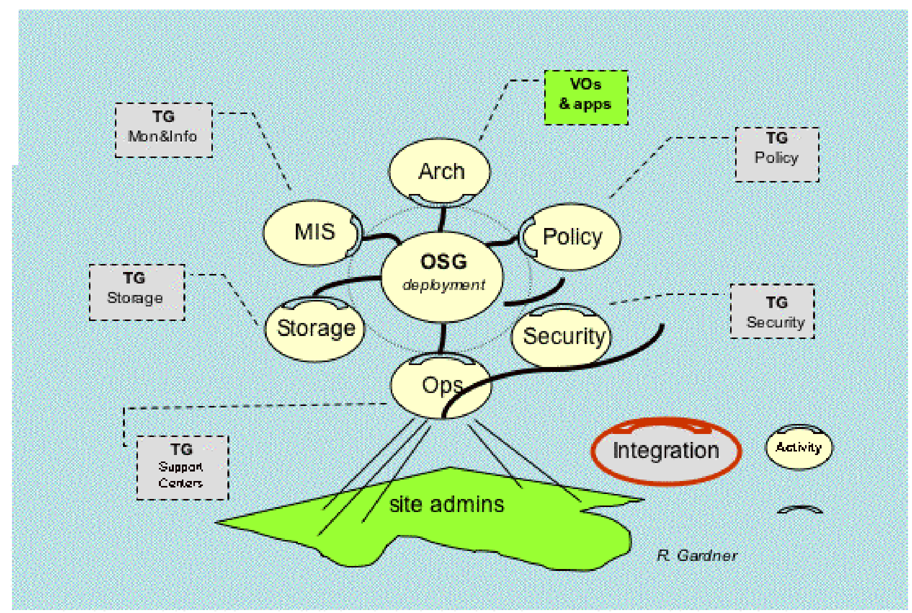

OSG Deployment Plan

Open Science Grid

- Evolve Grid3 to OSG in Spring 2005:
  - Flip the switch end of February.
  - Time-box of March and April to provision and consolidate.
- "Grid3Dev" iVDGL Grid Laboratory will integrate and validate new services.
- Joint projects contributing new and extended services:
  - Monitoring and Discovery infrastructure - University of Buffalo, University of Chicago, Caltech, US CMS, PPDG…
  - Storage Services - LBNL, US CMS, Fermilab, PPDG..
  - Account mapping and access control (AuthZ) - US ATLAS, US CMS, LCG, PPDG..
  - Operations - Indiana iGOC, iVDGL, LBNL, Fermilab..,

OSG Architecture

Open Science Grid

- OSG Blueprint documents principles and best practices to guide engineering, design and implementations:
  - The OSG architecture will follow the principles of symmetry and recursion.
  - Services should function and operate in the local environment when disconnected from the OSG environment.
  - Policy should be the main determinant of effective utilization of the resources.
- OSG promotes common interfaces in front of different implementations.
  - Sponsor testing and validation suites to support and ensure this.
  - Migration to WSRF & Web Services starting.
- No conceptual boundary between Grid wide and VO services.
  - OSG VO as "first class" entity.

Many Services must be added

Open Science Grid

- Storage resource access and management,
  - both to provide contracted persistent storage of data and management of data caches and temporary stores.
- Dataset management and caching,
  - meta-data services and management, wide area location and distribution of large scale data.
- Planning and optimization for effective use
  - discovery and scheduling
  - robust use of opportunistically available resources
- Multi-user access and support
  - from single, non-technical investigators to large cooperating groups within a managed organization.
- Diagnosis and troubleshooting
  - to manage the increase in scale and complexity.

Open Science Grid

Operations Is Key: Long list of responsibilities

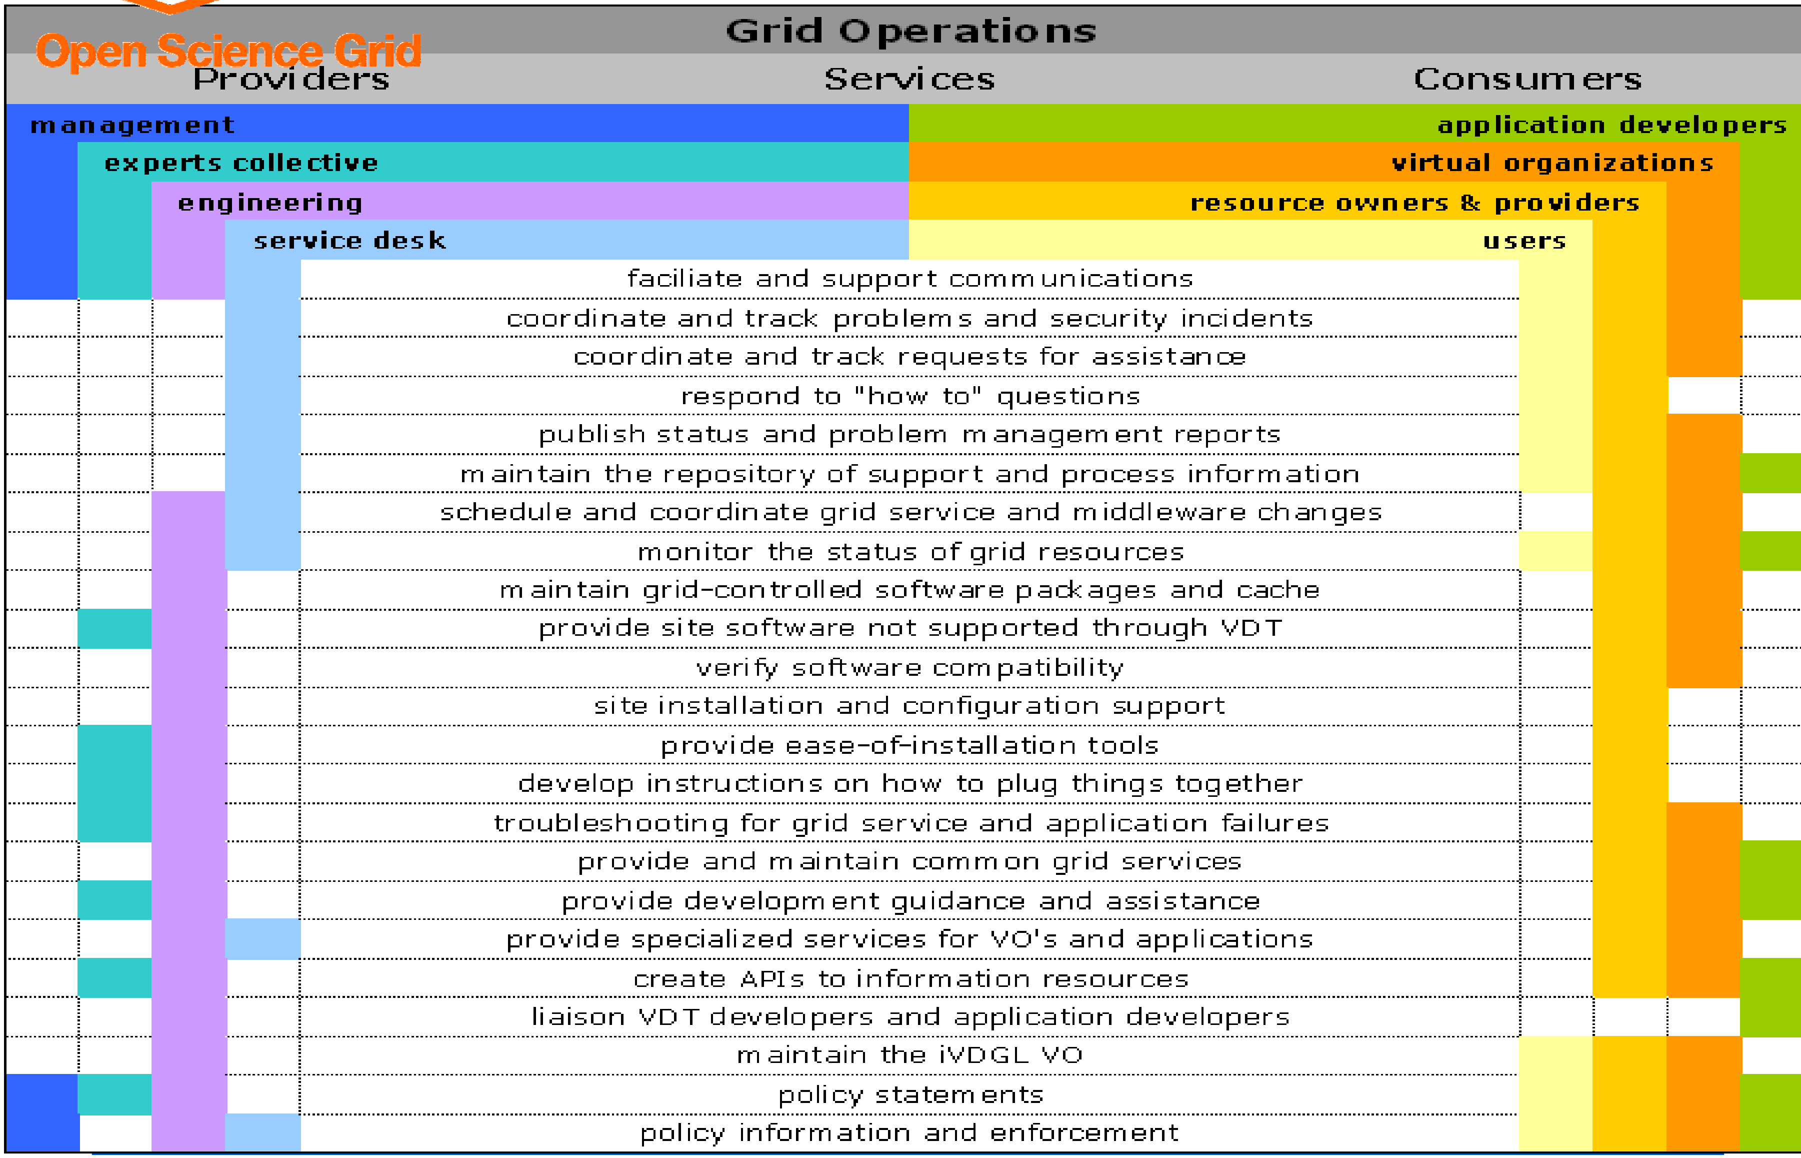

Open Science Grid

Challenges learned from Grid3

- Site & service providing perspective:
  - maintaining multiple “logical” grids with a given resource; maintaining robustness; long term management; dynamic reconfiguration; platforms
  - complex resource sharing policies (department, university, projects, collaborative), user roles
- Application developer perspective:
  - challenge of building integrated distributed systems
  - end-to-end debugging of jobs, understanding faults
  - common workload and data management systems developed separately for each VO

Opportunities facing OSG

Open Science Grid

- Build scalable, robust, effective set of Services.
- Achieve a common goal through community contributions.
- Use separate infrastructures as transparently accessible whole.
- Maintain operational commitment through decades long life-cycle of science community needs.

http://www.opensciencegrid.or
g